

REMARKS

This responds to the Office Action mailed on September 8, 2003 rejecting claims 1-46 and withdrawing claim 47 from examination. Claim 47 has been cancelled. Claim 48 has been added to recite the subject matter recited in claim 47, which is now cancelled. In view of the remarks set forth below, Applicant respectfully requests reconsideration, removal of the rejections, and allowance of all of the pending claims.

CLAIM REJECTIONS UNDER 35 U.S.C. 102(e)

CLAIM 1

Paragraph 4 of the Office Action rejects claim 1 under 35 U.S.C. 102(e) as being anticipated by Huitema (U.S. Patent No. 6,016,512). Paragraph 5 of the Office Action states that:

Huitema discloses a method for serving requests for Internet information files in an Internet caching system, comprising the steps of:

receiving, at a local Internet cache server, a user request from a user for an Internet information file [local cache server 310, Huitema Fig. 6, col. 3, lines 17-43];
in response to the received request, making a query for said information file, if said information file has not been cached by said local server [relay query, col. 1 line 63 - col. 2 line 7];

in response to a reply to said query, making a file request for said information file, wherein said, file request is directed to a feeder means [cache server 340, Fig. 6] if said reply indicates that a central file server [root server 340, Fig. 6; central cache server col. 3 lines 27-43], storing cached Internet information files, has said information file cached; and

querying, from said feeder means in response to said file request, said central file server for said information file, in order to decrease the load on said central file server [Huitema col. 5 lines 35-55, col. 6 lines 31-55]. It is clearly the cache server shared the load on central root server by periodically updating process [col. 4 line 52 – col. 5 line 13, col. 6 lines 47-55].

(Office Action, page 2, line 18-page 3, line 8)

Applicant traverses the rejection of claim 1.

Claim 1 recites a “method for serving requests for Internet information files in an Internet caching system, comprising the steps of: receiving, at a local Internet cache server, a user request from a user for an Internet information file; in response to the received request, making a query for said information file, if said information file has not been cached by said local server; in response to a reply to said query, making a file request for said information file, wherein said file request is directed to a feeder means if said reply indicates that a central file server, storing cached Internet information files, has said information file cached; and querying, from said feeder means in response to said file request, said central file server for said information file, in order to decrease the load on said central file server.”

As will be further discussed below, Huitema does not teach or suggest a method that includes “making a query for said information file, if said information file has not been cached by said local server; in response to a reply to said query, making a file request for said information file, wherein said file request is directed to a feeder means if said reply indicates that a central file server, storing cached Internet information files, has said information file cached; and querying, from said feeder means in response to said file request, said central file server for said information file”, as recited claim 1.

Huitema is directed to enhanced domain name servers and processing of domain name queries (col. 1, lines 14-16). Domain name servers are used to provide Internet Protocol (IP) addresses corresponding to domain names. For example, a user types in the domain name or universal resource locator (URL) of a remote computer using a browser at a local computer 110 (col. 1, lines 29-32). The local computer thereafter sends a query to a local server 120 to obtain the IP address of the remote computer (col. 1, lines 32-34). If the local server 120 does not find the answer in a local cache, the **local server 120 relays the query to a root server 130, which in turn transmits redirection information**, i.e., name and IP address of a remote server (col. 1, lines 39-46, emphasis added). **Thereafter, the local server 120 repeats the query, this time to the specified remote server**, which in turn transmits the IP address of the requested domain name (col. 1, lines 48-49, emphasis added). The local server 120 stores a copy of the answer and provides an answer to the local computer 110 (col. 1, lines 51-53). The local computer 110 uses the IP address to establish a connection to the remote computer 150 (col. 1, lines 54-55).

According to Huitema, the above domain name system has several problems. If the user mistypes or incorrectly guesses the domain name, the local server 120 will not find the domain name in the local cache and thus relays a query to the root server (col. 1, line 65-col. 2, line 2). The root server 130 thereafter checks a reference database, and if the domain name does not exist, the root server 130 sends an error message to the local server 120, which in turn relays the error message to the local computer (col. 2, lines 2-5). Thus, mistyping and wrong guesses can result in a transaction to one of the root servers, thus extending the domain name processing time and subjecting communications to problems with Internet traffic (col. 2, lines 12-15). Moreover, the caching technique no longer provides high quality results (col. 2, lines 16-18). Typically, fewer than 85% of queries are served in less than three seconds (col. 2, lines 18-19). This is expected to worsen as the Internet grows (col. 2, line 20).

To overcome the above problems, Huitema discloses an enhanced domain name system for more efficient processing of domain name queries (title, abstract, col. 1, lines 14-16). The system includes a local cache server 310, a cache server 340 and a root server 130 (col. 3, lines 19-22). The local cache server 310 is provided with a most frequently used (MFU) domain name table 320 and a validity code table 330 (col. 3, lines 22-24; col. 6, lines 47-49). This allows domain name queries to be **processed without accessing remote servers or root servers**, thereby reducing processing time (col. 6, lines 50-52, emphasis added). When a user types in the domain name of a remote computer, local cache server 310 accesses the MFU domain name table 320 (col. 5, lines 24-26). If local cache server 310 finds the answer, which will be the case in most instances, local cache server 310 transmits the corresponding IP address to local computer 110 (col. 5, lines 28-31). A different scenario occurs for domain name requests of invalid names (col. 5, lines 36-37). If the local cache server 310 does not find the answer, local cache server 310 generates a check code (col. 6, lines 34-38). If the generated check code is not in validity code table 330, local cache server 310 transmits an error message to local computer 110 (col. 6, lines 38-41). In the rare instance when the domain name is not in the MFU domain names table 320 and the check code matches one of the entries in validity code table 330, local cache server 310 processes the domain name query in the traditional manner (col. 6, lines 43-46). Unlike prior art systems, the domain name data in the local cache servers 310 is not collected by the local cache servers 310, but rather, by the network cache servers 340 (col. 3, lines 48-49). The network cache servers 340 provide the

local cache servers 310 with updates on a predetermined basis (col. 3, lines 51-54, emphasis added).

Thus, Huitema discloses a root server 130, a cache server 340 and a local cache server 310, which under certain circumstances, relays a query to the root server 130, which in turn transmits a message back to the local cache server 310. However, contrary to the assertion implied by the Office Action, Huitema does not teach or suggest making a request for information in response to the message from the root server 130, where such request is directed to the cache server 340. The only request made in response to the message from the root server 130 is directed to the address transmitted by the root server (i.e., the name and IP address of a remote server), not the cache server 340. Although the network cache server 340 provides the local cache server 310 with updates (e.g., new and revised domain name data), this is apparently done on a predetermined basis without a request for information directed to the cache server 340 (col. 3, lines 51-54).

Therefore, even if the cache server 340 constitutes a feeder, as asserted by the Office Action, and even if the root server 130 constitutes a central file server, as asserted by the Office Action, Huitema cannot possibly teach or suggest “making a query for said information file, if said information file has not been cached by said local server; in response to a reply to said query” and “making a file request for said information file, wherein said file request is directed to a feeder means if said reply indicates that a central file server, storing cached Internet information files, has said information file cached”, as recited in claim 1. Accordingly, Huitema also cannot teach or suggest “querying, from said feeder means in response to said file request, said central file server for said information file”, as recited in claim 1.

Consequently, Huitema does not teach or suggest a “method for serving requests for Internet information files in an Internet caching system, comprising the steps of: receiving, at a local Internet cache server, a user request from a user for an Internet information file; in response to the received request, making a query for said information file, if said information file has not been cached by said local server; in response to a reply to said query, making a file request for said information file, wherein said file request is directed to a feeder means if said reply indicates that a central file server, storing cached Internet information files, has said information file cached; and querying, from said feeder means in response to said file request, said central file server for said information file, in order to decrease the load on said central file server”, as recited in claim 1.

Accordingly, reconsideration and allowance of claim 1 is respectfully requested.

Claims 2-17 depend from claim 1 and are therefore patentable for at least the same reasons as stated above for claim 1. Accordingly, reconsideration and allowance of claims 2-17 is respectfully requested.

CLAIM 18

Paragraph 4 of the Office Action rejects claim 18 under 35 U.S.C. 102(e) as being anticipated by Huitema (U.S. Patent No. 6,016,512). Paragraph 12 of the Office Action states that:

As per claim 18, Huitema discloses an arrangement in an Internet caching system, said system comprising at least one local cache server and at least one central file server, both of which servers stores cached Internet information files, which arrangement, for decreasing the load on said central file server, includes a Feeder communicating with said local cache server and with said central file server [local cache server 310, cache server 340, DNS root server 130, Huitema Fig. 1, abstract], wherein said feeder includes:

first means for receiving a request for an Internet information file from said local cache server [Huitema col. 3 lines 27-43];

second means for deriving a query from an alphanumerical string received from said local cache server [hash coding, Huitema col. 5 lines 48 – col. 6 line 46]; and

third means for querying said central file server for said Internet information file using said query derived by said second means [corresponding name and address records, Huitema col. 3 lines 56 – col. 4 line 35].

(office action, page 4, line 10-page 5, line 3)

Applicant traverses the rejection of claim 18.

Claim 18 recites an “arrangement in an Internet caching system, said system comprising at least one local cache server and at least one central file server, both of which servers stores cached Internet information files, which arrangement, for decreasing the load on said central file server, includes a Feeder communicating with said local cache server and with said central file server, wherein said Feeder includes: first means for receiving a request for an Internet information file from said local cache server; second means for deriving a query from an alphanumerical string received from said local cache server; and third means for querying said central file server for said Internet information file using said query derived by said second means.”

Huitema does not teach or suggest an arrangement that includes a Feeder that includes: "first means for receiving a request for an Internet information file from said local cache server; second means for deriving a query from an alphanumerical string received from said local cache server; and third means for querying said central file server for said Internet information file using said query derived by said second means", as recited in claim 18.

As stated above, Huitema discloses a root server 130, a cache server 340 and a local cache server 310, which under certain circumstances, relays a query to the root server 130, which in turn transmits a message back to the local cache server 310. The local cache server 310 includes a list of the most active names together with the corresponding name and address records, i.e., (1) a record containing the domain name system names of domain name system servers that hold information about the specified domain and (2) a record containing the IP addresses of a computer identified by a domain name system name (col. 3, lines 56-61). The local cache server 310 also includes hash codes, or validity codes, which are preferably integer numbers, for use in error checking (col. 4, lines 13-18). However, contrary to the assertion implied by the Office Action, the name and address records are not derived from any query used in association with the hash coding. In addition, contrary to the assertion implied by the Office Action, the names and address records are not used for querying the root server 130.

Therefore, even if the cache server 340 constitutes a feeder, as asserted by the Office Action, and even if the root server 130 constitutes a central file server, as asserted by the Office Action, and even if the hash coding constitutes means for deriving a query from an alphanumeric string received from the local cache server 310, as asserted by the Office Action, Huitema cannot possibly teach or suggest "second means for deriving a query from an alphanumerical string received from said local cache server; and third means for querying said central file server for said Internet information file using said query derived by said second means", as recited in claim 18.

Consequently, Huitema does not teach or suggest an "arrangement in an Internet caching system, said system comprising at least one local cache server and at least one central file server, both of which servers stores cached Internet information files, which arrangement, for decreasing the load on said central file server, includes a Feeder communicating with said local cache server and with said central file server, wherein said Feeder includes: first means for receiving a request for an Internet information file from said local cache server; second means for deriving a query from an alphanumerical string received from said local cache server; and third means for querying

said central file server for said Internet information file using said query derived by said second means”, as recited in claim 18.

Accordingly, reconsideration and allowance of claim 18 is respectfully requested.

Claims 19-38 depend from claim 18 and are therefore patentable for at least the same reasons as stated above for claim 18. Accordingly, reconsideration and allowance of claims 19-38 is respectfully requested.

CLAIM REJECTIONS UNDER 35 U.S.C. 103

Paragraph 27 of the Office Action rejects claims 2-4, 22, 26, 28, 39-46 under 35 U.S.C. 103(a) as being unpatentable over Huitema (U.S. Patent No. 6,016,512) in view of Wessels et al (ICP and the Squid Web Cache). Paragraph 28 of the Office Action states that:

Huitema discloses a method a Internet caching system, comprising:

a set of local Internet cache servers, wherein each local cache server is arranged to receive request from users for Internet information files [local cache server 310, Huitema Fig. 6, col. 3 lines 17-43];

at least one central file server included in a central cache site and storing cached Internet information files [root server 340, Fig. 6; central cache server col. 3 lines 27-43];

feeder means interconnecting said set of local cache servers with said central file server [cache server 340, Fig. 6].

However Huitema did not teach means for communicating with at least one local cache server in accordance with a protocol used for communicating between Internet cache servers and means for retrieving Internet information files from said central file server using data base queries, thereby decreasing the load on said central file server.

It is well-known in the art that Internet Cache protocol (ICP) was used among Web caches to improve the exchange queries and replies [Wessels, abstract, page 345, col. 1 line 41 – col. 2 line 15], thereby reducing the workload from Web server.

Thus, it would have been obvious to one in the ordinary skill in the art at the time the invention was made to realize that using the ICP or Cache Digest would provide the cache knows whether or not the neighbor holds the requested data. Doing so would improve the data flow process between Web client nodes and servers over the large network.

(office action, page 7, line 16-page 8, line 12)

Applicant traverses the rejection on the grounds that the proposed combination of Huitema and Wessels is improper.

First, the Office Action has not presented legally sufficient evidence to support the proposed combination. The Examiner bears the burden of factually supporting any conclusion of obviousness, including evidence of a suggestion or motivation to combine the references and reasonable expectation of success of doing so. In order to try to satisfy this burden, the Office Action states that:

It is well-known in the art that Internet Cache protocol (ICP) was used among Web caches to improve the exchange queries and replies [Wessels, abstract, page 345, col. 1 line 41 – col. 2 line 15], thereby reducing the workload from Web server.

Thus, it would have been obvious to one in the ordinary skill in the art at the time the invention was made to realize that using the ICP or Cache Digest would provide the cache knows whether or not the neighbor holds the requested data. Doing so would improve the data flow process between Web client nodes and servers over the large network.

(Office Action page 8, lines 5-12)

However, this statement merely summarizes the role of the Internet Cache Protocol (ICP) as a mechanism for establishing complex cache hierarchies (see Wessels, page 346, col. 1, lines 28-35). It is not evidence that ICP or Squid would provide a benefit in the domain name system of Huitema. Moreover, although ICP may be of benefit in some types of systems, this does not automatically mean that it would be of benefit in the domain name system of Huitema. Yet the Office Action apparently boldly assumes that this is the case.

The combination is improper without proper evidence in support thereof. The mere fact that references can be combined or modified does not render the resultant combination obvious unless the prior art also suggests the desirability of the combination (MPEP 2143.01 citing *In re Mills*). In sum, the Office Action has not presented legally sufficient evidence to support the assertion that there is some teaching, suggestion, or motivation in the references themselves or in the knowledge available to one of ordinary skill in the art to make the proposed modification.

Rather, the Office Action has engaged in impermissible hindsight. In this instance, it is only Applicant's application which suggests the subject matter recited in the claims.

Second, Huitema does not teach or suggest any reason to modify the system to employ ICP or Squid. Moreover, even if one of ordinary skill in the art was looking to modify Huitema, why would they look to Wessels? Wessels isn't directed to processing domain name queries, as in Huitema. Rather, Wessels is directed to the Internet Cache Protocol (ICP) and its implementation in the Squid Web Caching software (abstract) for intermediate storage of copies of popular Web documents (page 345, col. 1, lines 39-44). Indeed, Wessels explicitly discusses the Domain Name Service (DNS), yet in discussing such, Wessels does not teach or suggest that ICP or Squid should (or could) be used in the Domain Name Service (Wessels, page 346, col. 1, lines 5-7).

Third, as discussed below, modifying the domain name system of Huitema to use ICP or Squid, as proposed by the Office Action, would merely extend processing time and increase network traffic without any apparent benefit to the system of Huitema. Thus, it would not have been obvious to combine Huitema and Wessels as proposed in the Office Action. One of ordinary skill in the art would not modify the system of Huitema if that modification would reduce the efficiency of the system without any apparent benefit. Accordingly, the proposed combination fails to meet the requirements of MPEP 2141.02, which states that a reference must be considered in its entirety, i.e., as a whole, including portions that would lead away from the claimed invention.

Huitema discloses an enhanced domain name system for storing domain names and efficiently processing domain name queries (title; col. 1, lines 14-16; col. 2, lines 28-29, emphasis added). The system **improves efficiency** by storing a most frequently used (MFU) domain name table and validity codes at local cache servers (col. 6, lines 47-49). (Note that ISP control cache servers 435 and edge servers 445 (FIG. 4) are examples of a local cache server 310 (see col. 3, lines 40-41)). This allows domain name queries to be **processed without accessing remote servers or root servers**, thereby reducing processing time (col. 6, lines 50-52). Unlike prior art systems, the domain name data in the local cache servers 310 is not collected by the local cache servers 310 themselves, but rather is provided by network cache servers 340 (col. 3, lines 48-49). The network cache servers 340 **provide** the local cache servers 310 with **updates** (e.g., new and revised domain name data) **on a predetermined basis** (col. 3, lines 51-54). For example, an update is necessary when a new domain name is added to central DNS database 410 (col. 5, lines

6-8). Cache administrator 425 selects updates or new information to propagate to ISP control cache servers 435 and edge cache servers 445 (col. 5, lines 3-6). Replicator 440 of ISP control cache server 435 replicates the updates or new information to the next level of cache servers, i.e., edge cache servers 445, which in turn replicates the updates or new information to other cache servers (col. 5, lines 9-13). Thus each of the cache servers appears to have the same information. In the rare instance when the domain name is not in the MFU domain name table 320, the local cache server 310 processes the domain name in the traditional manner (col. 6, lines 43-46).

Wessels is directed to the Internet Cache Protocol (ICP) and its implementation in the Squid Web Caching software (abstract, lines 1-3). ICP and Squid are demand driven caching systems (page 346, col. 2, lines 22-23) in which a cache receives a request for an object and thereafter sends a number of queries to peer caches asking whether they have a copy of such object (page 347, col. 2, lines 45-55). Each such peer cache receives one of such queries, checks for its existence locally and forwards a reply to the requesting cache (page 347, col. 2, lines 45-55). The requesting cache collects the replies and chooses a peer cache from which to retrieve the object (page 347, col. 2, final paragraph). After selecting a source for the object, the cache makes a regular HTTP request to retrieve it (page 348, col. 1, lines 7-8).

Modifying the domain name system of Huitema such that it uses the Internet Cache Protocol (ICP) or Squid Web Caching software, as proposed by the Office Action, would reduce the efficiency of the domain name system without any apparent benefit. As stated above, each of the cache servers in Huitema appear to have the same domain name information ("cache administrator 425 selects . . . updates or new information to propagate to ISP control cache servers 435 and edge cache servers 445 . . . Replicator 440 of ISP control cache server 435 replicates the updates or new information to the next level of cache servers, i.e., edge cache servers 445, which in turn replicates the updates or new information to other cache servers "; Huitema, col. 5, lines 3-13, emphasis added). As a result, in the rare instance that one of the caches does not have a requested domain name (col. 6, lines 43-46, emphasis added), there is no apparent incentive to query neighboring caches. That is, the neighboring caches will not have the domain name either. Indeed, sending queries to neighboring caches would merely extend processing time and increase network traffic without any apparent benefit. In view thereof, it would not have been obvious to combine Huitema and Wessels as proposed in the Office Action. One of ordinary skill in the art would not modify the system of Huitema if that modification would reduce the efficiency of the

system without any apparent benefit. Accordingly, the proposed combination fails to meet the requirements of MPEP 2141.02, which states that a reference must be considered in its entirety, i.e., as a whole, including portions that would lead away from the claimed invention.

Fourth, Wessels teaches that ICP has significant disadvantages. For example, there are numerous ways in which ICP fails to meet some of the demands placed upon it (page 351, col. 2, final three lines). An ICP query does not include any parent or sibling designation, so the receiver really has no indication of how the peer cache is configured to use it (page 352, col. 2, lines 1-3). Squid attempts to support this functionality with the miss_access feature, however, in addition to being awkward to implement, the miss_access feature brings its own complications (page 352, col. 2, lines 7-12). In addition, ICP still does not include any age information, neither in query nor reply (page 353, col. 1, lines 20-23). Thus, a cache may return an ICP_HIT if its copy of the object is fresh by its own configuration parameters, but the subsequent HTTP may result in a cache miss (page 353, col. 1, lines 22-25). Thus, situations emerge where the ICP reply no longer matches the HTTP request result (page 353, col. 1, lines 25-28). Another fundamental problem is that the ICP query does not provide enough information to accurately predict whether the HTTP request will be a hit or a miss (page 353, col. 1, lines 29-34). Still further, ICP has poor scaling characteristics (page 353, col. 32, lines 29-34) and a host of other disadvantages (pages 353-354).

Applicant has thus presented evidence that there is no suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the domain name system of Huitema to use ICP or Squid from Wessels. Applicant has also presented evidence to show that modifying the system of Huitema to use ICP or Squid, as proposed by the Office Action, would extend processing time and increase network traffic without any apparent benefit. Accordingly, the proposed combination fails to meet the requirements of MPEP 2141.02, which states that a reference must be considered in its entirety, i.e., as a whole, including portions that would lead away from the claimed invention. In view at least thereof, Applicant respectfully submits that it would not have been obvious to combine Huitema and Wessels as proposed in the Office Action. Thus, the rejection is improper.

Accordingly, reconsideration and withdrawal of the rejection is respectfully requested.

CONCLUSION

This application should now be in condition for allowance. A notice to this effect is respectfully requested. If the Examiner believes after this amendment, that the application is not in condition for allowance, the Examiner is requested to call the Applicants' attorney at the telephone number listed below.

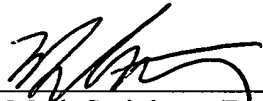
Because the reasons above are sufficient to traverse the rejections, Applicants have not explored, nor do they now present, other possible reasons for traversing such rejections. Nonetheless, Applicants expressly reserve the right to do so, if appropriate, in response to any future Office Actions.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicants hereby request any necessary extension of time.

If an additional fee is required, authorization is hereby given to charge such additional fees to Deposit Account No. 50-1402.

Respectfully submitted,

Date: February 9, 2004

By 
Mark Steinberg (Reg. No. 40,829)
MCCARTER & ENGLISH, LLP
City Place I
185 Asylum Street
Hartford, Connecticut 06103-3495
Ph.: (860) 275-7043
Fax: (860) 724-3397

HARTFORD: 608067.01
99999-00000